

Data Mountain: Using Spatial Memory for Document Management

*George Robertson, Mary Czerwinski, Kevin Larson,
Daniel C. Robbins, David Thiel, and Maarten van Dantzich*

Microsoft Research
One Microsoft Way
Redmond, WA 98052, USA
Tel: 1-425-703-1527
E-mail: ggr@microsoft.com

ABSTRACT

Effective management of documents on computers has been a central user interface problem for many years. One common approach involves using 2D spatial layouts of icons representing the documents, particularly for information workspace tasks. This approach takes advantage of human 2D spatial cognition. More recently, several 3D spatial layouts have engaged 3D spatial cognition capabilities. Some have attempted to use spatial memory in 3D virtual environments. However, there has been no proof to date that spatial memory works the same way in 3D virtual environments as it does in the real world. We describe a new technique for document management called the *Data Mountain*, which allows users to place documents at arbitrary positions on an inclined plane in a 3D desktop virtual environment using a simple 2D interaction technique. We discuss how the design evolved in response to user feedback. We also describe a user study that shows that the Data Mountain does take advantage of spatial memory. Our study shows that the Data Mountain has statistically reliable advantages over the Microsoft Internet Explorer Favorites mechanism for managing documents of interest in an information workspace.

KEYWORDS

3D user interfaces, desktop VR, information visualization, spatial cognition, spatial memory, document management

INTRODUCTION

Managing documents effectively on computers has been a key user interface design problem for the last thirty years. The issue has become more critical as users venture onto the World-Wide Web, because the number of easily accessible documents has increased dramatically. Graphics technology, processor speed, and primary memory capacity advances have made it possible to build systems that

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

UIST '98, San Francisco, CA

© 1998 ACM 0-58113-034-1/98/11... \$5.00




Figure 1: Data Mountain with 100 web pages.

help with this document management problem.

The Data Mountain (Figure 1) is a novel user interface for document management designed specifically to take advantage of human spatial memory (i.e., the ability to remember where you put something). In our current prototype, the user freely arranges document thumbnails on an inclined plane textured with passive landmarks. We use 3D visual and audio cues to enhance the similarity to real-world object arrangement, yet use simple 2D interaction techniques and common pointing devices (like the mouse) for all interactions. The system is designed with a fixed viewpoint, so users need not navigate around the space. Users can identify and distinguish documents both through their thumbnail representation and through pop-up titles.

In this paper we describe the document management task and discuss existing graphical solutions. We then discuss related work from the field of spatial cognition, and issues of navigation and document management in the specific context of the World-Wide Web. Next, we describe the Data Mountain in detail, and report on a user study that compared it to the Microsoft Internet Explorer (IE4) Favorites mechanism. The paper concludes with a discussion of study findings and planned future work.

Document Management

Document management tasks occur in a variety of contexts, and over a wide range of sizes of information stores and information structures. For example, tasks include managing files in a file system, mail messages, and web pages. The basic information structures a user might encounter include unordered sets, ordered lists, hierarchies, and graphs. Often the documents may belong to more than one of these information structures. Document browsing, searching, overviews, histories, and information workspaces all employ such structures.

The concept of information workspaces, introduced by Card, Robertson, and Mackinlay in 1991 [4], refers to the environment in which documents of interest can be compared, manipulated, and stored for future use. Iconic desktops, web browser Favorites or Bookmarks, and the Web Forager [5] are examples of information workspaces.

RELATED WORK

Document Management Systems

Early graphical methods for document management included list views, expandable lists for viewing hierarchies, and iconic 2D spatial layouts. The Apple Macintosh (circa 1984) included list views and a spatial layout (icon view). The spatial layout allowed the user to place icons in whatever grouping the user desired. Apple later added expandable lists for hierarchies, and piles [18]. Piles enrich the spatial layout by allowing the user to group related documents and take less screen space for the group. Yet another form of 2D spatial layout is the Treemap [14], which is a space-filling layout that is generated automatically, used primarily for overviews of document collections and their meta-data.

SemNet [10] was an early 3D spatial layout of documents. It tackled the difficult problem of visualizing networks. The result was difficult to understand because of the complexity of the information and layout. The Information Visualizer project [4][22] at Xerox PARC introduced a broad set of 3D visualization and interaction techniques for understanding information. In 1994, Maya Design Group introduced Workspace [3] as the first example of a 3D spatial layout of documents under the user's control. The Web Forager [5] built on the experience gained from the Information Visualizer project and introduced a 3D spatial layout for web pages and WebBooks [5].

The Visible Language Workshop of the MIT Media Lab did research in the design of dynamic virtual information spaces, combining typography and 3D graphics layout to present visually appealing interactive information landscapes [16][21].

Typical 2D desktops (Windows or Mac OS) use a spatial layout of icons and overlapping windows. As will be seen, the Data Mountain supports a larger number of objects, prevents overlap with a page avoidance algorithm, and uses thumbnail images instead of icons.

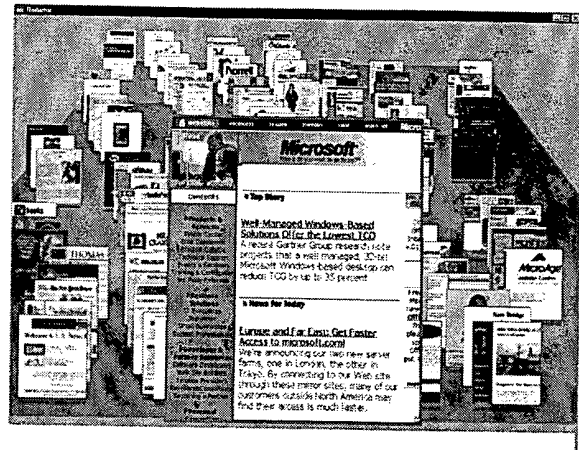


Figure 2: Selected page in preferred viewing position.

PadPrints [13] uses the Pad++ 2D zoomable user interface to implement a thumbnail image-based web history mechanism that is superior to text based history mechanisms. The authors raise a question about whether their technique is successful because of the thumbnail images or because of the zoomable interface. PadPrints uses an automatic layout for short term use; the Data Mountain uses a manual layout to exploit spatial memory for long term use.

Many of the techniques mentioned made use of spatial cognition, whether or not this was done intentionally. In particular, automatic spatial layouts of information leverage the user's ability to recognize and understand spatial relationships (both in 2D and 3D). The 3D interfaces make it possible to display more information without incurring additional cognitive load, because of pre-attentive processing of perspective views (i.e., smaller size indicates spatial relationships at a distance). The Maya Design Group Workspace system and the Web Forager intended to use spatial memory, by allowing the user to place documents as an aid to finding them again.

The Data Mountain is an advance over Workspace and Web Forager in several ways. First, the Data Mountain allows the user to place the document at an arbitrary location on its slope with a simpler interaction technique than the earlier systems (taking advantage of constrained motion along the inclined plane of the Mountain). Second, when a page is being moved, other pages are moved out of the way (active page avoidance), yet the user still sees visual cues indicating where every page will be when the movement is completed. Third, the Data Mountain exploits a variety of audio cues to augment the visual cues. Fourth, page titles are displayed whenever the mouse moves over a page. And fifth, in light of research in spatial cognition and wayfinding, visual neighborhood demarcation cues are provided to assist the user in arranging her personal space on the Data Mountain.




Figure 3: Second subject's layout of same 100 pages.




Figure 4: Third subject's layout of same 100 pages.

Spatial Cognition

There is a large body of literature on spatial cognition (see [23][7] for recent examples) and wayfinding [8][9][25], both for real and electronic worlds. Some of these studies have culminated in a set of guidelines for designers of virtual worlds. For instance, leveraging knowledge from the architectural domain [17][19], Darken and Silbert [9] have shown that adding real world landmarks, like borders, paths, boundaries and directional cues, can greatly benefit navigation performance in virtual reality. In their studies, Darken and Silbert have shown that stationary or predictably moving cues are optimal, and that multiple sensory modalities can be combined to assist searching through an electronic space (like 3D sound cues). They also have shown that if the space is not divided using a simple, organizing principle, users will impose their own, conceptual organization upon the space.

DATA MOUNTAIN

The Data Mountain is a 3D document management system. The current prototype is being used as an alternative to current web browser Favorites or Bookmark mechanisms, so we sometimes refer to pages as the objects that appear on a mountain. It should be understood that other forms of documents should work equally well.

When a page is first encountered, it appears in a preferred viewing position (see Figure 2), so that it is easily read. The user can place the page by dragging it with a traditional left-mouse-button drag technique. As the page is being dragged, other pages move out of the way so the page being moved is not occluded. After the page has been placed, it can be selected with a single click to bring it back to the preferred viewing position. Visual and audio cues, as well as the interaction techniques are more fully described below.

The Data Mountain is designed to work in a desktop 3D graphics virtual environment (also known as desktop VR),

although it could certainly work in either Fishtank VR [27] or VR with head-tracked head-mounted displays. Examples described here are desktop VR examples.

Leveraging Natural Human Capabilities

The primary motivation for the design of the Data Mountain came from a desire to leverage natural human capabilities, particularly cognitive and perceptual skills. In particular, 3D perception is used to allow for the representation of a large number of web page thumbnails with minimal cognitive load. Our pre-attentive ability to recognize spatial relationships based on simple 3D depth cues (like perspective views and occlusion) makes it possible to place pages at a distance (thereby using less screen space) and understand their spatial relationships without thinking about it. We can leverage audio perception to reinforce what is happening in the visual channel. Both the visual and auditory perception can enable basic human pattern recognition capabilities. And finally, we hope to use spatial memory to make it easier to find documents in the information workspace.

In the real world, spatial memory often aids us in finding things. For example, when we place a piece of paper on a pile in our office, we are likely to remember approximately where that paper is for a long time. Our hope is that this ability works as well in a virtual space as it does in the physical world. This is not an obvious conclusion. However, if spatial memory is primarily an act of building a mental map of the space, then we should be able to do the same thing in a virtual environment, and take advantage of it.

Data Mountain Visual and Audio Design

The Data Mountain provides a continuous surface on which documents are dragged. The document being dragged remains visible so the user is always aware of the surrounding pages. This is in direct contrast to the way a Web Forager [5] user places documents in a discrete set of tiered locations using flicking gestures. We believe the

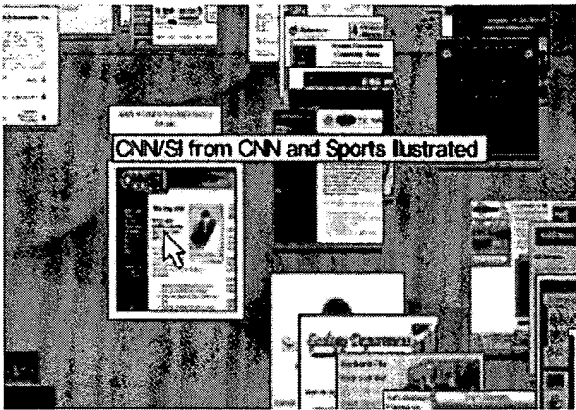


Figure 5: Title shown while hovering over page.

user's act of directly placing the page on the continuous surface of the Data Mountain aids spatial memory.

The Data Mountain prototype uses a planar surface (a plane tilted at 65 degrees), as shown in Figure 1. The landscape texture on the Data Mountain surface provides passive landmarks for the user meant as an aid for grouping objects into categories, but the landmarks have no explicit meaning. The user can place the web pages anywhere on the mountain. In practice, users create meaning by organizing the space. In our study, there were many ways to lay out the same set of pages (compare Figures 1, 3, and 4). There is no "right" layout; rather, the layout is very personal, has meaning for the individual who created it, and evolves over time under user control.

Note that the current prototype provides no mechanism for labeling groups of pages with category titles. While some users requested this feature, we found they built a very accurate mental map of their categories even without explicit labels. Some users employed particularly salient thumbnails as visual identifiers of their groups, keeping them in front of all other members of the group (thus creating their own landmarks).

There are a number of 3D depth cues designed to facilitate spatial cognition. The most obvious are the perspective view and occlusion, particularly when pages are being moved. The landmarks also offer an obvious cue, which may or may not be utilized during page placement as well as retrieval. Less obvious, but also quite important, are the shadows cast by the web pages.

Subtle but pervasive audio cues accompany all animations and user actions to reinforce the visual cues. The sound effects are highly dynamic. For example while moving a page the user hears a humming sound that changes pitch based on the speed of the page as it is dragged, as well as indicating spatial location by controlling volume, low pass filtering, panning, and reverb level. As the user moves a page, other pages move out of the way as needed, producing another distinctive sound.

Data Mountain Interaction Design

When the user clicks on a page stored on the Data Mountain, the page is moved forward to a preferred viewing position, as shown in Figure 2. The animation to bring the page forward lasts about one second [4], uses a slow-in/slow-out animation [6][11], and is accompanied by an audio cue. We use a higher resolution texture map for the page image in the preferred viewing position, ensuring that the page is quite readable.

When in the preferred viewing position, a click on the page will either select and follow a hyperlink, or put the page back on the Data Mountain in its last known location. This is also done with a one second, slow-in/slow-out animation accompanied by audio.

We provided a pop-up label similar to tool-tips to display page titles. Subjects tended to use their spatial memory to get to the neighborhood of the page, then riffle through the titles (like riffling through a pile of papers on your desk) to find the page. A standard tool-tip uses a hover time before the tip is displayed. We determined in a pilot study that the hover time was not effective since it precluded rapid inspection of multiple titles. Hence, the title appears as soon as the mouse moves over a page. An example is shown in Figure 5. In one group from our user study, the title was shown just above (but disconnected from) the page. Some subjects could not easily distinguish which thumbnail the title applied to, so our second Data Mountain design added an identically colored halo around the thumbnail, creating a visual link to the title.

A page can be moved at any time by dragging it with the mouse. Since the page is visible during the move, the user knows where the page will be when the drag is terminated. The movement is continuous and constrained to the surface of the Data Mountain. This results in one of the principal advantages of the Data Mountain; the user gets the advantages of a 3D environment (better use of space, spatial relationships perceived at low cognitive overhead, etc.), but interacts with it using a simple 2D interaction technique.

Page Avoidance Behavior

When moving a page, what is the right behavior for pages that are encountered (i.e., how are collisions handled)? We have tried three alternatives, each improvement driven by user comments. First, we did nothing. That is, the page in motion simply passed through the pages in the way. This approach suffers from several problems. It makes the system seem lifeless and makes the metaphor harder to understand. In addition, it is quite easy to put pages right on top of each other, making it difficult to find some pages.

Second, we tried a simulation of tall grass. Think about what happens when you walk past tall grass. If you walk slowly, the grass moves out of your way slowly then returns. If you walk fast, the grass seems to fly out of your way. We implemented a simple simulation in which previously placed pages behave like grass displaced with a

page the user is dragging. Each page that is moved out of the way uses a 1/2 second animation to move aside, followed by a one second animation to move back. This feels very lively but suffers from two problems. It still does not solve the problem that two pages can end up in the same location: the user may drop a page close to another document's real location while that document is temporarily displaced, soon to return, causing one to occlude the other. The collision avoidance behavior should be designed to eliminate such surprises. Also, the movement is based on first encounter. In other words, if after triggering avoidance the dragged page slows down or hovers in-place, the return animation will still take place, causing the objects to intersect anyway. Essentially, the 'grassy technique' works well for continuous dragging but tends to be annoying when you slow down. Note that slowing down is an essential part of object placement: effectively, the grassy technique caused some users to create unnecessary occlusion since they would pick the just vacated spot of a displaced page for the new page. It could be modified to be dependent on your speed; but since it is basically an estimate, it will fail in some cases. The first group of Data Mountain users described in the experiment below used this page avoidance mechanism. These problems provided strong motivation for finding another method and running another set of subjects, as some users effectively lost many pages due to occlusion.

In our current implementation, we continually maintain a minimum distance between all pages, even while a page is being moved, and transitively propagate displacement to neighbors as necessary. This has the advantage that the user dragging the page continually sees what state will result when the drag is terminated (i.e., there is no animation settling time). Also, the pages never get fully obscured. In particular, you cannot move a page and leave two pages in the same location. On the other hand, displacements may propagate far afield when a cluster of closely packed pages is 'pushed' by a dragged page, resulting in more visual unrest than is really desirable. Still, the approach feels quite lively, and was used by the final group of subjects in our study. We feel this change contributed most to improved user performance in the second Data Mountain group.

Implementation

The Data Mountain prototype runs under Windows NT version 4 on PCs equipped with Intergraph Intense 3D Pro 1000 or Pro 2200 graphics accelerators. All application code was written in C++, utilizing our own libraries for animation and scene-graph management. These libraries in turn used the ReActor [2] infrastructure and OpenGL as the underlying graphics library.

The interactive sound for the Data Mountain was based on MISS (the Microsoft Interactive Sound Sequencer) which takes parametric sound events and sequences them using MIDI to communicate to the wavetable synthesizer on a Creative Labs AWE64 Gold card.

For prototyping and user study purposes, the web pages in the current implementation are not "live", i.e.; one cannot select and follow a hyperlink. Future versions of the Data Mountain will contain live web browsing capabilities.

The 100 pages used in the study below are screen snapshots of actual web pages in 24-bit color. We employ two bitmap sizes of each page for texture mapping: a small 64x64 pixel version (12KB each) for the thumbnails on the Data Mountain surface, and a 512x512 pixel version (768 KB each) for the close-up view. One hundred thumbnails plus a close-up together will fit in 2MB of texture cache. Our system implements text labels using texture-mapped fonts [12]. This is vastly preferable to vector fonts, and has the advantage that it enables display of legible text on surfaces that are not screen-aligned. Category labels placed on the Data Mountain surface are just one potential use of such perspective-distorted text.

If pop-up labels are naively attached to thumbnails, they will be subject to perspective projection, and thus be smaller for pages that are placed towards the back of the information workspace. We found it difficult to choose an absolute label size in model coordinates that produces appropriately scaled text for labels of both foreground and background documents. Instead, we implemented labels for thumbnails to be of identical size, independent of the document's distance from the viewer in virtual space. We do this by placing the title-tip a constant distance away from the eye-point, on the vector from the eye-point to the page whose title we are showing.

USER STUDY

Studies by Tauscher and Greenberg [24] and Abrams [1] were the earliest attempts to gather information about user behavior as they traverse the web over several months. According to Abrams [1], users develop their own personal web information spaces through the use of Favorites mechanisms in order to combat the problems of information overload, pollution, entropy, structure and lack of a global view of the web. Users do this by building a smaller, more valuable, organized and personal view of the web.

Usage tracking shows that hotlists, bookmarks and Favorites folders are the navigation tools most frequently utilized by users for locating information on the web [20].

Hence, web browser designers need to provide their users with mechanisms for creating personal web information spaces that can reliably and efficiently return the user to their favorite web sites. Implementing such mechanisms relaxes the cognitive and temporal demands of hypertext navigation [1]. Usability studies, as well as basic research, however, indicate that the current designs for navigating the web are still sub-optimal in supporting users' cognitive models of web spaces and the amount of information they need to repeatedly consume [1][24].

The Abrams study [1], in particular, pointed out how episodic memory [26], or memory for events, could be

thought of as a primary cognitive avenue for retrieving web pages from Favorites. As pointed out in that research, reviewing a set of bookmarks or Favorites is basically a process of using textual cues to retrieve a memory for how that web page was stored. Usage data from that study showed that users had trouble retrieving their favorite web pages, often because the default title which they used to store the web page was an inadequate retrieval cue for recognition. Left unanswered, though, was whether and to what extent spatial cognition was playing a role in users' navigational behaviors, and to what extent our browser designs could leverage what is known about spatial cognition and wayfinding. As users' personal web information spaces grow larger, how can we effectively design the Favorites user interface to afford efficient retrieval? One way to increase the amount of information on the screen is to move to a 3D user interface. Also, the possibility of adding spatial landmarks, edges and audio cues, as discussed above, should prove useful. These are the issues that are explored in the current set of studies.

The primary user interface design for a Favorites folder is similar to the hierarchical tree views used to browse files in a computer file system. A user is allowed to enter items into an organized list, often alphabetized, and the list can have any number of subcategory structures added to it. Often, the URL or the web page title is the default label representing the web page when browsing the list. The list is text only, so does not allow users to leverage other channels of information that may also be effective when attempting to retrieve web pages, such as the auditory channel. The question of interest is how effective is the Data Mountain for leveraging all aspects of memory during the retrieval of a web page in Favorites?

Our hypothesis was that the effectiveness and usability of the Data Mountain depends in part on the transfer of real-world spatial memory skills to a virtual environment.

This assumption was also made for earlier systems that make use of spatial memory for document management (like the Maya Designs Workscape and the Xerox PARC Web Forager). To test this hypothesis, we did a user study comparing use of the Microsoft Internet Explorer 4.0 (IE4) Favorites and the Data Mountain for the same storage and retrieval tasks. IE4 was chosen for this comparison because it is a shipping product with which many readers will be familiar.

Methods

Subjects. Thirty-two experienced IE4 users participated in this study. All users had to successfully answer a series of screening questions pertaining to web browser and Internet knowledge in order to qualify for participation. Subject ages ranged from 18 through 50 years old, and all had normal or corrected-to-normal vision. The number of females and males was balanced.

Equipment. The study was run on high-end Pentium machines (P6-266 or P6-300), with at least 128 MB of memory, and a 17-inch display. The machines had either an Intergraph Intense 3D Pro 1000 or 2200 graphics accelerator card and ran Windows NT4. One hundred web pages were used in this study; fifty pages were selected randomly from PC Magazine's list of top web sites and fifty pages selected randomly from the Yahoo! database. A web server was contained on the local computer for each user eliminating network lag. This was done to maintain consistent system response times across applications, users and time of study, thus eliminating possible confounds in the results.

Procedure. Users participated in one of three groups. One group of users stored and retrieved web pages using the standard IE4 Favorites mechanism. The other two groups of users stored and retrieved web pages using the Data Mountain. The second Data Mountain group (DM2) interacted with a version of the Data Mountain that incor-



Title	Title, Summary, & Thumbnail	Summary	Thumbnail
Bezerk - The Free On-line Entertainment Network	Bezerk - The Free On-line Entertainment Network The new, FREE premier online entertainment network features YOU DON'T KNOW JACK the netshow, the online version of the irreverent quiz show party game CD-ROM! 	The new, FREE premier online entertainment network features YOU DON'T KNOW JACK the netshow, the online version of the irreverent quiz show party game CD-ROM!	

Table 1: Examples of the four cueing conditions used in the study




Figure 6: Average retrieval times for each condition.




Figure 7: Average number of incorrect pages retrieved by cueing condition and application.




Figure 8: Average number of failed trials by cue condition and application.

porated design changes suggested from the first Data Mountain group (DM1). The design changes included preventing pages from occluding each other during storage or manipulation, strengthening the association between "hover titles" and their corresponding pages in a more intuitive manner, and improving the audio feedback (e.g. adding spatialization effects). The discussion section (below) will detail to what extent the design changes for the second Data Mountain group improved user satisfaction and performance.

Participants in each group were told that they would be storing 100 web pages. They were allowed to create any organizational structure they wanted in the IE4 Favorites mechanism or the Data Mountain, and were encouraged to create a structure that mimicked how they stored favorite web pages at home or work. Their web page organizations from home were collected for comparison purposes. The participants were told they would have to use their organization for tasks in the second half of the test session, though were not told what the second half of

the session would entail. The IE4 users could use either the Favorites panel or the Organize Favorites dialog box for storage and organization. After all pages had been stored, participants in each group were given an additional opportunity to re-organize their Favorites. Note that there are differing interaction constraints on the users depending on the application they used. When using IE4, users were restricted in terms of the default amount of the screen space taken up by the Favorites mechanism. The Data Mountain users, on the other hand, used the entire screen for managing their web page layout by default. However, since Data Mountain users were operating on cropped images of web pages, not the "live" pages, they were restricted in terms of how much they could learn about a web page (e.g., they could not scroll down and read the contents of a page for more information like the IE4 users could and often did). All study participants were experienced IE4 users.

During retrieval, which ensued after a short break, the participants were shown one of four different retrieval cues and asked to find the related page. The four retrieval cueing conditions were: the title of the page, a one or two sentence summary of the page's content, a thumbnail image of the page, and all three cues simultaneously (called the "All" cue). Participants saw 25 trials of each cueing condition, for a total of 100 retrievals. All pages presented for retrieval were seen in the earlier storage phase. The web pages to be stored and the subsequent retrieval cues were presented in a random order for each participant. Table 1 shows an example of each of the four styles of retrieval cues. If a participant could not find the target page within two minutes, a "time-out" was enacted and the participant was instructed to proceed to the next retrieval task. Page Retrieval was defined as selecting an item from the IE4 Favorites to be displayed in the main browser pane, or bringing a page forward using the Data Mountain. Users were not explicitly discouraged from producing incorrect retrievals.

Four main dependent variables were used in this study: (1) web page retrieval time; (2) the number of incorrect pages selected prior to finding the correct page; (3) the number of trials for which the participant failed to retrieve the correct page within the two-minute deadline; and (4) the participants' subjective ratings of the software. These dependent measures are assumed to be powerful indices of subjects' abilities to locate items in space. If they know where their web pages are, retrieval performance should be efficient.

RESULTS

Reaction Time. The main finding in the reaction time data was that the Data Mountain reliably facilitated speedy retrieval of web pages when compared to IE4, allowing users to leverage visual as well as textual cues in finding document locations in 3D space. Also, the two applications supported different retrieval behaviors. IE4 leveraged the available textual, title information—any other or additional kind of information that was presented as a retrieval cue had a deleterious effect on performance. For both groups of Data Mountain users, the All cue enabled users to retrieve pages faster than using just a title, indicating that the Data Mountain prototype lets users utilize additional information modalities for improved spatial location memory. Figure 6 shows the results. Figures 6-8 include error bars showing plus or minus one standard error from the mean. Statistical tests that support these findings are presented next.

A 3x4 (Application x Cue Condition) analysis of variance (ANOVA) with repeated measures was performed on the reaction time data. The analysis revealed a statistically reliable main effect of application, $F(2,18) = 4.84$, $p < .02$, a statistically reliable effect of cueing condition, $F(3,27) = 5.7$, $p < .01$ and a statistically reliable interaction between application and cueing condition, $F(6,54) = 3.32$, $p < .01$. Post hoc analyses (Scheffe tests) revealed that in IE4 the title was the retrieval cue that resulted in

the fastest reaction times, reliably faster than either the thumbnail or summary cues (but not reliably faster than the All cue, though this was close to being reliably different). In the two Data Mountain groups the pattern of results were very different. Participants were reliably faster, on average, using the Data Mountain, especially in the thumbnail and All cueing conditions. The only condition in which the first Data Mountain group was slower than the IE4 group was the title cueing condition. The second Data Mountain group was as fast or faster than the first Data Mountain group and IE4 group in all cueing conditions.

Number of Incorrect Retrievals. In essence, users performed more accurately with the second Data Mountain. These results are shown in Figure 7. Statistical analyses supporting these findings are presented next.

A 3x4 (Application x Cue Condition) ANOVA with repeated measures revealed a statistically reliable main effect of application for number of incorrect pages retrieved, $F(2,18) = 4.48$, $p < .03$, as well as a reliable effect of cue, $F(3,27) = 7.62$, $p < .001$. There was also a reliable interaction between application and cueing condition, $F(6,54) = 2.4$, $p < .04$. For the IE4 group incorrect pages were most often visited when the cue was a thumbnail or a summary, while the two Data Mountain groups were only reliably more likely to visit an incorrect page when a summary cue was provided. Reliably fewer incorrect pages were retrieved in the second Data Mountain group than the other two groups.

Failed Trials. The two Data Mountain groups were reliably more likely to retrieve a web page within the time limit than the IE4 group. The IE4 group was more likely to fail trials in either the summary or thumbnail condition, while the Data Mountain groups were more likely to fail trials in the summary condition. This data is shown in Figure 8, with statistical support presented below.

A 3x4 (Application x Cue Condition) ANOVA with re-

Questionnaire Item	IE4	First Data Mountain	Second Data Mountain
I like the software.	3.4 (1.0)	3.3 (1.2)	3.7 (0.7)
The software is efficient.	3.6 (1.1)	2.9 (1.2)	3.3 (0.8)
The software is easy to use.	3.6 (1.1)	4.0 (0.9)	4.0 (1.1)
The software feels familiar.	4.0 (0.7)	3.3 (1.2)	3.4 (1.2)
It is easy to find the page I am looking for with the software.	3.4 (1.0)	3.3 (1.0)	3.4 (1.0)
Organizing web pages is easy with the software.	3.4 (1.1)	3.7 (1.1)	3.8 (0.8)
If I came back a month from now I would be still be able to find many of these web pages.	3.2 (0.8)	4.1 (0.6)	3.6 (1.3)
I was satisfied with my organization scheme.	3.2 (1.2)	3.4 (0.8)	3.1 (1.2)
My organizing scheme was very similar to the organization in my home Favorites folder.	3.9 (1.5)	3.6 (1.4)	3.4 (1.6)

Table 2: User satisfaction averages for on a five point scale where 1=disagree, 5=agree (standard deviations in parentheses)

peated measures revealed reliable main effects of application for the number of trials where the user could not retrieve the page before the two minute deadline, $F(2,18) = 8.3$, $p < .01$, as well as a reliable effect of cueing condition, $F(3,27) = 46.9$, $p < .001$, and the interaction between application and cue condition, $F(6,53) = 19.51$, $p < .001$.

Subjective Ratings. After completing the retrieval tasks, the participants answered questions about their satisfaction with the application they used in the study. Table 2 shows the average and standard deviation scores on a five-point scale (1=disagree, 5=agree) for participants' responses to a number of ease of use and likability ratings. A one-way analysis of variance was run on each measure to test for differences between the groups. No reliable main effect was found for any of the ratings.

One final question was presented for both Data Mountain groups, inquiring whether they would prefer to use IE4 or the Data Mountain software. There was a reliable preference for IE4 in the first Data Mountain group, and eight out of eleven of the second Data Mountain group said they would prefer to use the Data Mountain over IE4 (one participant in the second Data Mountain group failed to answer this question). This is a statistically reliable preference for the second Data Mountain using a binomial test, and clearly provides converging evidence, beyond the performance data, that iterative testing with end users has improved the user interface of the Data Mountain.

DISCUSSION AND FUTURE WORK

The user study reported here demonstrates that the Data Mountain is an effective alternative for current web Favorites mechanisms, even in this preliminary prototype form. The Data Mountain allows users to informally arrange their space in a very personal way. This informality appears to have great power, and is enabled by having the ability to view the whole space and the spatial relationships between the pages, as well as manually control those relationships in space.

The user study also suggests that spatial memory does in fact play a role in 3D virtual environments. We often heard subjects say things like "it's right here", or "I know it's back there", and move directly to the location of the page. Storage times, retrieval times, and retrieval failures were all reduced because of this aspect of spatial memory's influence. However, previous research [15] has suggested that little significant value is provided by adding spatial location information to the storage and subsequent retrieval of a document over and above simply providing a semantic label for the same purposes. In that study, users were able to accurately and efficiently retrieve stored documents in the real world with as impoverished a semantic label as a two-letter cue! Storing the document in a spatial position did improve performance over baseline control conditions, however. We intend to continue our research as to how these real world results

map onto the retrieval of information in large, unstructured, electronic worlds, using alternative measures, perhaps more indicative of spatial cognition's contributions to performance. We also intend to investigate these issues over a longer period of retrieval time.

There are a number of additional areas we want to explore further. We would like to understand the relative contributions to this successful study of the various components (3D versus 2D, spatial memory, audio, title display, page avoidance, thumbnail images). It is clear from the last study group that the right page avoidance mechanism can make a big difference. As with PadPrints, it is possible that the thumbnail images are a significant contributor.

The Data Mountain was first conceived as a contoured landscape with various landmarks on it. The only constraint is that the contours need to be monotonically increasing (i.e., there should be no valleys in which documents could be hidden from view). Based on our experience with the planar version, it appears that careful choice of the landscape texture placed on the surface could afford most, if not all, of the advantages of a contoured surface. Nevertheless a contoured Data Mountain landscape should be experimentally compared with the planar surface on the current prototype. Also, a number of subjects asked about how hard it would be to make the passive texture landmarks become manipulable (to move groups of pages, for example). In addition, we have begun to experiment with implicit grouping, based on proximity and "white space" between groups. It may be a great benefit to the user to have a page being moved reflect through some visual property the implicit group that it is currently nearest.

Finally, we have begun exploration of camera-based head motion parallax to enhance the usability of the Data Mountain. Preliminary results look very promising.

ACKNOWLEDGEMENTS

We would like to acknowledge Susan Dumais, Matt Conway, and Ken Hinckley for their helpful advice during all stages of this research.

REFERENCES

1. Abrams, D. (1997). Human Factors of Personal Web Information Spaces, M.S. Thesis, University of Toronto, 1997, <http://virtual.inesc.pt/rct/rct.51.html>.
2. Ball, J.E., Ling, D.T., Pugh, D., Skelly, T., Stankowsky, A., & Thiel, D. (1994) ReActor: A System for Real-Time Reactive Animations. In *CHI'94 Conference Companion* (Boston, MA, April 24-28, 1994), ACM press, 39-40.
3. Ballay, J.M. (1994). Designing Workspace: An Interdisciplinary Experience. In *CHI'94 Conference Proceedings* (Boston, MA, April 24-28, 1994), ACM press, 10-15.

4. Card, S.K., Robertson, G.G., & Mackinlay, J.D. (1991). The Information Visualizer, an Information Workspace. In *Proceedings of CHI '91, Human Factors in Computing Systems* (New Orleans, LA, April 27-May 2, 1991), ACM press, 181-188
5. Card, S.K., Robertson, G.G., & York, W. (1996). The WebBook and the WebForager: An Information Workspace for the World-Wide Web. In *Proceedings of CHI '96 Human Factors in Computing Systems* (Vancouver, BC, Canada, April 13-18, 1996), ACM press, 111-117.
6. Chang, B. & Ungar, D. (1993). Animation: From Cartoons to the User Interface. In *Proceedings of UIST '93* (Atlanta, GA, November 3-5, 1993), ACM / SIGGRAPH, SIGCHI, NY, 1993, 45-55.
7. Curiel, J.M. & Radvansky, G.A. (1998). Mental organization of maps. In *Journal of Experimental Psychology: Learning, Memory and Cognition*, 24, 202-214.
8. Darken, R. & Sibert, J. L. (1993). A Toolset for Navigation in Virtual Environments. In *Proceedings of ACM User Interface Software & Technology (UIST)*. New York: ACM.
9. Darken, R. & Sibert, J. L. (1996). Navigating large virtual spaces. *International Journal of Human-Computer Interaction*, 8, 49-72.
10. Fairchild, K. M., Poltrock, S. E. and Furnas, G. W. (1988). Semnet: three-dimensional graphic representations of large knowledge bases. In *Cognitive science and its applications for human-computer interaction*, Guindon, R. (ed), Lawrence Erlbaum, 1988.
11. Gonzalez, C. (1996). Does Animation in User Interfaces Improve Decision Making? In *Proceedings of CHI '96 Human Factors in Computing Systems* (Vancouver, BC, Canada, April 13-18, 1996), ACM press, 27-34
12. Haeberli, P., and Segal, M. (1993). Texture Mapping as a Fundamental Drawing Primitive. In *Proceedings of the Fourth Eurographics Workshop on Rendering*, Michael Cohen, Claude Puech, Francois Sillion, Editors. Paris, France, June, 1993.
13. Hightower, R.R., Ring, L.T., Helfman, J.I., Bederson, B.B., and Hollan, J.D. (1998). Graphical Multiscale Web Histories: A Study of PadPrints. In *Proceedings of ACM Conference on Hypertext 1998*.
14. Johnson, B. and Shneiderman, B. (1991). Space-filling approach to the visualization of hierarchical information structures. In *Proceedings IEEE Visualization '91*, pp. 284-291.
15. Jones, W. & Dumais, S. (1986). The spatial metaphor for user interfaces: Experimental tests of reference by location versus name. *ACM Transactions of Office Information Systems*, 4, pp. 42-63.
16. Kullberg, R.L. (1995). Dynamic Timelines: Visualizing Historical Information in Three Dimensions. *M.S. Thesis at the MIT Media Laboratory* (Boston, MA, September, 1995)
17. Lynch, K. (1960). *The Image of the City*. Cambridge, Massachusetts: The MIT Press.
18. Mander, R., Salomon, G. & Wong, Y.Y. (1992). A 'Pile' Metaphor for Supporting Casual Organization of Information. *CHI '92*, 627-634.
19. Passini, R. (1984). *Wayfinding in Architecture*. New York: Van Nostrand Reinhold.
20. Pitkow, J., and Recker, M. (1996). "The GVU 5th WWW user survey" http://www.cc.gatech.edu/gvu/user_surveys/
21. Rennison, E. & Strausfeld, L. (1995). The Millenium Project: Constructing a Dynamic 3+D Virtual Environment for Exploring Geographically, Temporally and Categorically organized Historical Information. In: *Spatial Information Theory: A Theoretical Basis for GIS*, International Conference COSIT '95, (Semmering, Austria, September 21-23, 1995), Lecture Notes in Computer Science, Vol. 988, Springer, 69-91
22. Robertson, G.G., Card, S.K., & Mackinlay, J.D. (1993). Information Visualization Using 3D Interactive Animation. *Communications of the ACM*, 36(4), April '93, 56-72.
23. Roskos-Ewoldsen, B., McNamara, T.P., Shelton, A.L. & Carr, W. (1998). Mental Representations of Large and Small Spatial Layouts are Orientation Dependent. In *Journal of Experimental Psychology: Learning, Memory and Cognition*, 24, 215-226.
24. Tauscher, L., & Greenberg, S. (1997). Revisitation Patterns in World Wide Web Navigation. *ACM SIGCHI'97 Conference on Human Factors in Computing Systems*. Atlanta, Georgia, March 22-27, ACM Press.
25. Thorndyke, P. W. & Hayes-Roth, B. (1982). Differences in Spatial Knowledge Acquired from Maps and Navigation, *Cognitive Psychology*, 14, pp. 560-589.
26. Tulving, E. & Thomson, D. M. (1973). Encoding Specificity and Retrieval Processes in Episodic Memory. *Psychological Review*, 80, 342-373.
27. Ware, C., Arthur, K., and Booth, K. (1993). Fish Tank Virtual Reality, *CHI'93 Proceedings*, 37-42.